

A More Sensitive Context

Christopher J. Banks and Ian Stark

Laboratory for Foundations of Computer Science
School of Informatics, University of Edinburgh, UK

Abstract. *Logic of Behaviour in Context (\mathcal{LBC})* is a spatio-temporal logic for expressing properties of continuous-state processes, such as biochemical reaction networks. \mathcal{LBC} builds on the existing *Metric Interval Temporal Logic* (MITL) and adds a “context modality” that explores the behaviour of a system when composed with an external process. \mathcal{LBC} models are terms of the *Continuous π -Calculus* ($c\pi$), a process algebra with continuous state space.

Our previously published \mathcal{LBC} model-checking technique required examining many points along the behavioural trajectory of a process; and potentially computing further trajectories branching off at every such point. This raised two difficulties: mixing temporal and spatial modalities could require computing a large number of trajectories, with costly numerical solution of differential equations; and might still fail to check intermediate values between discrete points on those trajectories.

In this paper we make progress against both of these problems using techniques from signal temporal logic and from sensitivity analysis. Boolean signals aggressively compress trace information, allowing more efficient computation; and sensitivity analysis lets us reliably check formulae over a region by calculating a smaller number of sample trajectories.

1 Introduction

The Logic of Behaviour in Context (\mathcal{LBC}) [3] is a spatio-temporal logic for expressing temporal and contextual properties of continuous state processes (dynamical systems), such as biochemical reaction processes. Temporal properties express behaviour over time and contextual properties express behaviour in the presence of another process. \mathcal{LBC} equips a metric interval temporal logic with a *context modality* which asserts the behaviour of a process when composed with another process. As such, the logic requires a model with a suitable notion of process composition; in this case we use the Continuous π -Calculus ($c\pi$) [11]. We give a brief account of $c\pi$ in Section 2 which is sufficient for the reader to understand its use as a model for \mathcal{LBC} and in Section 3 we give the syntax and semantics of the logic.

In previous work [3] we gave various model checking algorithms that can be used to verify the satisfaction of an \mathcal{LBC} formula over a $c\pi$ model. Each of these algorithms takes an approximate approach to model-checking continuous dynamical systems, using discrete simulation traces computed by numerical

solvers. This is a non-exact approach, but well founded in practice for the efficient approximate model checking of continuous processes (e.g. Antonioti et al. [2], Fages and Rizk [9], or Nickovic and Maler [14]). The approach relies on the assumption that the numerical solver is suitably precise, and that the discrete sample points of the simulation trace are sufficiently dense to give a good approximation of the dynamics of the system.

Our method for model checking of the context modality relied on the same assumption. To compute the satisfaction of a context modality we take sample points along the trajectory of the model, compose with the new context (process) at each point, and verify that the desired behaviour is observed on trajectories computed from these. A major problem with this approach is that to be sure the sample points are dense enough for a good approximation one has to numerically solve a large number of new trajectories, which is computationally costly. This approach also has the problem that there is no way to ensure the property tested is not violated by trajectories starting between the discrete sample points.

In this paper we address both of these problems: using Boolean signals to compress traces, following Maler and Nickovic [12]; and sensitivity analysis to reduce the number of sample trajectories, building on Donzé and Maler [8].

2 Continuous π -calculus

The continuous π -calculus ($c\pi$) was designed as a formal language for the study of evolutionary variation in biochemical processes. The canonical reference for $c\pi$ is Kwiatkowski's thesis [10], but the original language semantics were first published by Kwiatkowski and Stark [11].

The language syntax is based on the π -calculus of Milner [13], with some alterations and additions to better support the description of biochemical models. The usefulness of π -calculus style languages for modelling biochemical processes is well established, having first been described by Regev et al. [16,15].

The description of a biochemical process in $c\pi$ is split into two levels: *species* and *process*. A *species* in $c\pi$ is a description of the behaviour and interaction capability of a biochemical species. This level is similar to a π -calculus term. A *process* in $c\pi$ is a real-indexed parallel composition of each of the species in the biochemical process, representing a mixture with some initial concentration of each species.

Kwiatkowski and Stark [11] give full details, but for this paper it is only necessary to understand that a $c\pi$ process is a parallel composition of species S , each with an initial concentration c . The species in a composition may, or may not, interact with each other.

Definition 1 (Process). *The set \mathcal{P} of $c\pi$ processes is defined by the following grammar:*

$$(Process) P, Q ::= c \cdot S \mid P \parallel Q$$

A process may be a species S , as defined in [11], with initial concentration $c \in \mathbb{R}$, or a composition (mixture) of species.

A $c\pi$ process has a compositional semantics in terms of real vector spaces, from which an initial value problem, which gives the trajectory of a process, can be extracted. This is solved normally by a numerical simulator which outputs a time series for each of the species in a process. The space within which we place these systems and the trajectories of their behaviour over time is *process space*.

Definition 2 (Process space). *The process space \mathbb{P} is the vector space $\mathbb{R}^{(\mathcal{S}^\#)}$, where $\mathcal{S}^\#$ is the set of prime species — elementary species which cannot be broken down into a composition of two non-trivial species.*

3 Logic of Behaviour in Context

The Logic of Behaviour in Context (\mathcal{LBC}) [3] arose from the desire to define a logic for $c\pi$ which would allow the classification of the behaviour of a $c\pi$ model. It was clear that a temporal logic, and a logic which allowed the expression of constraints on the real-valued concentrations of $c\pi$ species, was required.

However, particularly in biochemical systems, behaviour is often reasoned about in terms of, not just the system itself, but the system’s behaviour when it is perturbed somehow. We wish to be able to reason about the system’s behaviour in some external context. Thus \mathcal{LBC} was conceived.

\mathcal{LBC} combines $\text{LTL}(\mathbb{R})$ (see Calzone et al. [5]) and MITL (see Alur et al. [1]) with the addition of a *context modality*. Here $\text{LTL}(\mathbb{R})$ is a temporal logic for properties of the real-valued concentrations of biochemical species; and MITL adds concrete times and time intervals to those temporal modalities.

The context modality ($Q \triangleright \psi$) holds for a process P whenever ψ holds in the presence of a new process Q . This allows the expression of behaviour in some given context. For example, consider the assertion $(c \cdot In) \triangleright \mathbf{G}([Pr] < x)$. In a biochemical context this could represent “in the presence of a concentration c of inhibitor In the concentration of product Pr in the system always remains below x ”.

The context modality is based on the guarantee operator from Cardelli and Gordon’s spatial logic [6]. However, that guarantee takes an arbitrary formula on the left hand side ($\phi \triangleright \psi$) to give a formula that holds for processes that when combined with any other process satisfying ϕ give a combination satisfying ψ :

$$P \models \phi \triangleright \psi \iff \forall Q (Q \models \phi \implies P \parallel Q \models \psi).$$

Model-checking a logic with this guarantee is hard because of the necessity to quantify over all processes that satisfy an arbitrary formula. Caires and Lozes [4] give an account of the undecidability of spatial logic with the guarantee.

The context modality, however, gives some of the power of guarantee in a more computationally feasible form by specialising the left hand side to a specific process:

$$P \models Q \triangleright \psi \iff Q \parallel P \models \psi$$

Nicola and Loreti [7] take a similar approach with their logic **MoMo** whose a production operator is based on guarantee. However, **MoMo** is defined for mobile

processes with resources and locations, modelled using a formalism based on shared tuple spaces, and the semantics of the production operator is incompatible with the kind of continuous state processes we address here.

3.1 Syntax

The syntax of \mathcal{LBC} follows MITL [1], with propositional atoms being inequalities between arithmetic combinations of real-valued concentrations of species in the system and their time derivatives. To this we add the context modality $Q \triangleright \phi$.

Definition 3 (\mathcal{LBC} formula). *The set Φ of \mathcal{LBC} formulae ϕ, ψ is defined by the following grammar:*

$$\begin{aligned} \phi, \psi &::= Atom \mid \phi \wedge \psi \mid \phi \vee \psi \mid \phi \Rightarrow \psi \mid \neg \phi \\ &\mid \phi \mathbf{U}_I \psi \mid \mathbf{F}_I \phi \mid \mathbf{G}_I \phi \mid Q \triangleright \phi \\ Atom &::= True \mid False \mid Val \bowtie Val \\ Val &::= v \in \mathbb{R} \mid [S] \mid [S]' \mid Val \oplus Val \\ \bowtie &::= > \mid < \mid \geq \mid \leq \\ \oplus &::= + \mid - \mid \times \mid \div \end{aligned}$$

where relational operators \bowtie and arithmetic operators \oplus have the standard meaning, $[S]$ denotes the concentration of species S , $[S]'$ denotes the rate of change over time of the concentration of species S , Q is a $c\pi$ process, and $I \subseteq \mathbb{R}^+$ is any non-negative time interval.

We use the abbreviations \mathbf{U} , \mathbf{F} , and \mathbf{G} to denote $\mathbf{U}_{[0, \infty]}$, $\mathbf{F}_{[0, \infty]}$, and $\mathbf{G}_{[0, \infty]}$ respectively. Likewise, for $t \in \mathbb{R}_{\geq 0}$, we write \mathbf{U}_t , \mathbf{F}_t , and \mathbf{G}_t as abbreviations for $\mathbf{U}_{[0, t]}$, $\mathbf{F}_{[0, t]}$, and $\mathbf{G}_{[0, t]}$ respectively.

3.2 Semantics

We define the semantics of \mathcal{LBC} by its satisfaction relation \models , where $P \models \phi$ if and only if $c\pi$ process P satisfies formula ϕ .

Definition 4 (Atomic propositions of a $c\pi$ process). *The set $\text{Props}(P)$ of atomic propositions satisfied by a process P is defined by*

$$\begin{aligned} True &\in \text{Props}(P) \\ False &\notin \text{Props}(P) \\ v_1 \bowtie v_2 \in \text{Props}(P) &\iff \text{value}(v_1, P) \bowtie \text{value}(v_2, P) \end{aligned}$$

where the relational operators \bowtie are defined in the normal way and

$$\begin{aligned} \text{value}(v, P) &= v \text{ for } v \in \mathbb{R} \\ \text{value}([S_i], P) &= c_i \text{ where } P = c_1 \cdot S_1 \parallel \dots \parallel c_n \cdot S_n \\ \text{value}([S_i]', P) &= c'_i \text{ where } dP/dt = c'_1 \cdot S_1 \parallel \dots \parallel c'_n \cdot S_n \\ \text{value}(v_1 \oplus v_2, P) &= \text{value}(v_1, P) \oplus \text{value}(v_2, P) \end{aligned}$$

where arithmetic operations \oplus are similarly defined as normal.

Definition 5 (\mathcal{LBC} satisfaction relation). For $P \in \mathcal{P}$, a $c\pi$ process, and \mathcal{LBC} formulae ϕ and ψ the satisfaction relation \models is defined inductively as follows:

$$\begin{aligned}
P \models \text{Atom} &\iff \text{Atom} \in \text{Props}(P) \\
P \models \phi \wedge \psi &\iff P \models \phi \text{ and } P \models \psi \\
P \models \neg\phi &\iff P \not\models \phi \\
P \models \phi \mathbf{U}_I \psi &\iff \text{for some } t \in I, P^t \models \psi \text{ and for all } t' \in [0, t], P^{t'} \models \phi \\
P \models Q \triangleright \phi &\iff (Q \parallel P) \models \phi
\end{aligned}$$

where $Q \in \mathcal{P}$ is any $c\pi$ process and process P^t is the state reached by process P after time t ; that is, the concentration of each species in P^t will be those present after P has run for time t . The notation P^t is shorthand for a function mapping $\mathcal{P} \times \mathbb{R}^+ \rightarrow \mathcal{P}$, with $P^0 = P$.

The remaining propositional connectives are derived as usual, together with temporal modalities $\mathbf{F}_I \phi \equiv \text{True} \mathbf{U}_I \phi$ and $\mathbf{G}_I \phi \equiv \neg(\mathbf{F}_I(\neg\phi))$.

Definition 6 (Duration of formula). The duration of a formula — the length of time to which it refers — is defined inductively:

$$\begin{aligned}
|\text{Atom}| &= 0 & |\phi \wedge \psi| &= \max(|\phi|, |\psi|) & |\neg\phi| &= |\phi| \\
|Q \triangleright \phi| &= 0 & |\phi \mathbf{U}_{[a,b]} \psi| &= \max(|\phi|, |\psi|) + b .
\end{aligned}$$

4 Signal checking

In our earlier paper [3] we used a model-checking method based on direct computation over time series from numerical simulation. Here we refine this technique using Boolean signals, inspired by Maler and Nickovic [12].

The key idea for signal-based temporal model checking is that the dynamics of the model is represented as a set of Boolean signals. Basic signals represent whether an atomic proposition is satisfied at a given time; for each logical operator we have a matching combinator for signals; and by applying these recursively over the structure of a formula we compute its satisfaction signal.

Simple formula satisfaction for a process is computed by taking the initial value of the corresponding signal: $P \models \phi \iff (s_\phi(0) = \text{True})$

Signals have a compact representation as sets of time intervals, in general giving significant compression over time-series of data values. The signal combinators can be efficiently implemented for this representation, which leads to substantial performance improvements over our earlier \mathcal{LBC} implementation [3].

4.1 Signals

Signals are constructed from the dynamics of a model. A signal represents the satisfaction of a formula at any given time. The set *Signal* is the set of finite length Boolean signals.

Definition 7 (Finite length Boolean signal). A finite length Boolean signal s of length r is a function $s : [0, r) \rightarrow \mathbb{B}$. A finite length Boolean signal has finite variability and, therefore, may be represented by a finite interval covering. For a signal s with length r an interval covering is a sequence $\mathcal{I} = I_1, I_2, \dots$ of left-closed right-open intervals such that $\bigcup I_i = \mathcal{I}$ and $I_i \cap I_j = \emptyset$ for all $i \neq j$. The minimal covering \mathcal{I}_s of the signal s is consistent with the signal if $s(t) = s(t')$ for all t, t' in the same interval $I_i \in \mathcal{I}_s$. A covering \mathcal{I}' is a refinement of \mathcal{I} , denoted $\mathcal{I}' \prec \mathcal{I}$, if for all $I' \in \mathcal{I}'$ there exists $I \in \mathcal{I}$ such that $I' \subseteq I$. The set of positive intervals of s is $\mathcal{I}_s^+ = \{I \in \mathcal{I}_s : s(I) = \text{True}\}$ and the set of negative intervals is $\mathcal{I}_s^- = \mathcal{I}_s \setminus \mathcal{I}_s^+$.

Signal checking relies on the conversion from the dynamics of the model to a set of *basic signals*, each of which represent the satisfaction of an atomic proposition.

Basic signals Basic signals are constructed from simulation traces of the form $[(t_0, \mathbf{c}_0), \dots, (t_n, \mathbf{c}_n)] \in \text{Trace}$, where each t_i is a time point and \mathbf{c}_i is a vector of the species concentrations at that time. To construct the basic signals we use the following procedure:

1. Take each leaf ϕ in the syntax tree of the formula; these are the atomic propositions of the form $[A] \bowtie c$, *True*, or *False*.
2. For each ϕ we construct a signal s_ϕ as an interval covering \mathcal{I} of intervals $[t_0, t_1), [t_1, t_2), \dots, [t_{n-1}, t_n)$.
3. Each interval $[t_i, t_{i+1})$ is in $\mathcal{I}_{s_\phi}^+$ if the constraint in ϕ is satisfied by the values in \mathbf{c}_i , otherwise it is in $\mathcal{I}_{s_\phi}^-$.

This set of signals captures the value over time of all the atomic propositions from a formula.

Signal combinators For non-atomic formulae we have a set of signal combinators that take the basic signals, apply a logical operation, and give the signal for the satisfaction of a formula over time. A signal s is constructed by computing its covering intervals \mathcal{I}_s ; it is sufficient to compute the positive intervals \mathcal{I}_s^+ as the negative intervals \mathcal{I}_s^- are, by definition, complimentary.

Definition 8 (Boolean signal combinators [12]). The signal combinators apply the logical connectives (\neg, \wedge) and temporal modalities (\mathbf{F}, \mathbf{U}) to signals (s_ϕ, s_ψ) and are defined as follows:

$\neg s_\phi$

Negation is a simple negation of the signal such that $\mathcal{I}_{s_{\neg\phi}}^+ = \mathcal{I}_{s_\phi}^-$.

$s_\phi \wedge s_\psi$

For conjunction we first compute a refinement of the coverings $\mathcal{I}_\phi^R \prec \mathcal{I}_\phi$ and $\mathcal{I}_\psi^R \prec \mathcal{I}_\psi$ such that $\mathcal{I}_\phi^R = \mathcal{I}_\psi^R$ and is the sequence of intervals I_1^R, \dots, I_n^R . The conjunction is then computed interval-wise such that $s_{\phi \wedge \psi} = s_\phi \wedge s_\psi$. The minimal covering $\mathcal{I}_{s_{\phi \wedge \psi}}$ is then computed by merging any adjacent intervals of the same Boolean value.

$\mathbf{F}_{[a,b]} s_\phi$

The temporal $\mathbf{F}_{[a,b]}$ modality is computed by back-shifting the positive intervals. $\mathcal{I}_{\mathbf{F}_{[a,b]} \phi}^+$ is constructed by taking each interval $I \in \mathcal{I}_\phi^+$ and computing its back-shifting $I \ominus [a,b] \cap \mathbb{R}^+$ where $[m,n] \ominus [a,b] = [m-b, n-a]$ and the intersection with \mathbb{R}^+ eliminates any negative times. The minimal covering $\mathcal{I}_{\mathbf{F}_{[a,b]} \phi}$ is then computed by merging any adjacent intervals of the same Boolean value.

$s_\phi \mathbf{U}_{[a,b]} s_\psi$

The fundamental temporal $\mathbf{U}_{[a,b]}$ modality can be computed on the basis that $\phi \mathbf{U}_{[a,b]} \psi \iff \phi \wedge \mathbf{F}_{[a,b]}(\phi \wedge \psi)$ when s_ϕ is a unitary signal. A signal s is unitary if \mathcal{I}_s^+ is a singleton. So if s_ϕ is unitary and it holds at t_1 and t_2 then it must hold for the whole interval $[t_1, t_2]$. For the case where s_ϕ is not unitary we can decompose it into a set of unitary signals $\{s_\phi^1, \dots, s_\phi^n\}$ and compute, for each $i \in [1, n]$:

$$s_\phi^i \mathbf{U}_{[a,b]} \psi = s_\phi^i \wedge \mathbf{F}_{[a,b]}(s_\phi^i \wedge s_\psi)$$

The signal is then recomposed to give:

$$s_\phi \mathbf{U}_{[a,b]} \psi = \bigvee_{i=1}^n s_\phi^i \mathbf{U}_{[a,b]} \psi$$

Context modality signal Our problem is how to compute the signal for the context modality $Q \triangleright \phi$. The positive intervals of the signal $\mathcal{I}_{s_{Q \triangleright \phi}}^+$ must represent the times at which if Q is composed with the model then ϕ is satisfied. To compute this we must choose a finite number of arbitrary time points at which to introduce Q to the model and compute the satisfaction of ϕ . The problem lies in how to choose these time points.

An initial solution is to choose the same time points as in the original trace; that is if we are checking $P \models q \triangleright \phi$ then we use the same time points as in the trace for P . The assumption here is that if the chosen time points for P were sufficiently dense then they will be sufficiently dense for $Q \triangleright \phi$.

Definition 9 (Context modality signal). *The context modality signal is constructed as follows. To compute a signal for $P \models Q \triangleright \phi$, for each time point t in the originally computed simulation trace we compute a new process $P^t \parallel Q$. Each of these new processes is solved numerically to get a trace and we recursively apply the signal checking procedure to find whether or not ϕ holds for each of these processes. The Boolean result from each process at time t is the value of the signal for the interval $[t, t')$ where t' is the time of the next point.*

The use of signal checking has, in a preliminary implementation, shown good improvements in performance for checking the temporal fragment of the logic. However, the focus of this paper is to show that improvements can also be made in the following technical issues.

4.2 Outstanding issues

There are two main issues with this approach to checking $P \models Q \triangleright \phi$. If we wish to check a temporal nested context modality, say $P \models \mathbf{G}_{[0,t]}(Q \triangleright \phi)$, then taking sample points on the original trajectory of P and checking $P^{t'} \parallel Q \models \phi$ for each sample point $t' \in [0, t]$ does not necessarily ensure that $P^{t''} \parallel Q \models \phi$ holds for t'' between sample points. We have simply worked on the assumption that with sufficiently dense sample points we will have great enough coverage to have a reasonable degree of confidence in the result. This seems a reasonable assumption to make, given that this applies to all applications of numerical simulation.

The second issue is that having a very large set of sample points means that we have to make a very large number of calls to the solver in order to solve $P^t \parallel Q$ for each sample point t . Experience with real examples has shown this is easily the most significant factor in the computational cost of model-checking \mathcal{LBC} .

5 Context checking with sensitivity

Our new method addresses both of the problems outlined in Section 4.2. The method potentially reduces calls to the ODE solver whilst also improving the coverage of context introduction between the original sample points. The key to the method comes from a study of sensitivity analysis for safety properties by Donzé and Maler [8]. Sensitivity analysis is used to systematically check a system with uncertain initial conditions; the use of sensitivity analysis ensures that a few discrete simulations cover a continuous space of initial conditions. Here we can apply the same principle allowing $P^t \parallel Q \models \phi$ to be checked for all t , because the difference between $P \parallel Q$ and $P^t \parallel Q$ is just a change in initial conditions. This initial condition space is potentially large, but we can find a sufficient and finite set of samples using an adaptation of Donzé and Maler’s technique; thus potentially reducing the number of calls made to the ODE solver.

5.1 Expansion function

If we take the set X_0 to be the set of initial conditions for the dynamical system corresponding to the processes $P^{t_0} \parallel Q, \dots, P^{t_n} \parallel Q$, where Q is introduced to P at times in $[t_0, t_n]$, then we take $\xi_{x_0}(t)$ to be the trajectory which is the unique solution for some initial condition $x_0 \in X_0$. The set of states reachable within time t from any initial condition is denoted $\mathbf{reach}_{\leq t}(X_0)$ and the set of states reachable at exactly time t is denoted $\mathbf{reach}_{=t}(X_0)$. Let d be a distance metric between points and let it extend to points and sets thus $d(x, X) = \inf_{y \in X} (d(x, y))$.

The set $\mathcal{B}_\delta(x)$ is the δ -ball around point x and the set $\mathcal{B}_\delta(X)$ is the δ -ball around set X thus $\mathcal{B}_\delta(X) = \bigcup_{x \in X} \mathcal{B}_\delta(x)$. Donzé and Maler show that it is possible to find the ball which tightly over-approximates $\mathbf{reach}_{=t}(\mathcal{B}_\delta(x_0))$ by means of the *expansion function*. Therefore it is possible to construct a “flow tube” around a trajectory which tightly over-approximates the reachable set.

Definition 10 (Expansion function [8]). Given $x_0 \in X_0$ and some $\epsilon > 0$ the expansion function of trajectory ξ_{x_0} , denoted by $\Xi_{x_0, \epsilon} : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ maps t to the smallest non-negative number δ such that all trajectories with initial state in $\mathcal{B}_\epsilon(x_0)$ reach a point in $\mathcal{B}_\delta(\xi_{x_0}(t))$ at time t :

$$\Xi_{x_0, \epsilon}(t) = \sup_{d(x_0, x) \leq \epsilon} d(\xi_{x_0}(t), \xi_x(t))$$

The value of the expansion function is the radius of the tightest ball around the reachable set from the ϵ -ball around the initial condition. The key here is that if we take the initial set to be the ball which tightly bounds our possible initial conditions then we can compute an over-approximation of the reachable set and therefore prove that we do not reach a state where some ϕ holds. Donzé and Maler [8] show that $\Xi_{x_0, \epsilon}(t)$ can be computed via *sensitivity to initial conditions* $\frac{\partial \xi_{x_0}}{\partial x_0}(t)$, which is commonly implemented by numerical solvers. The error in the numerical approximation is quadratic in ϵ [8]. Therefore, for our algorithm we have a parameter θ which is the maximum initial ball radius; if the ball around our initial conditions is greater than θ then we refine until the initial balls are all smaller than θ .

Donzé and Maler then have a scheme for refining the over-approximation for an arbitrary set of initial conditions in \mathbb{R}^n . However here our initial conditions are less general and so, in Section 5.2, we give a more specific method used for computing the signal of a context modality.

5.2 Application to checking \mathcal{LBC}

We can apply the principle of computing traces of flow tubes, where necessary, instead of traces of trajectories. Some preliminary definitions follow.

The set B is the set of balls in \mathbb{P} and for $\beta \in B$: $c(\beta)$ is the centre point of the ball and $r(\beta)$ is the radius. The set $Tube$ is the set of flow tube traces in \mathbb{P} where a flow tube trace is of the form $(t_0, \beta_0), \dots, (t_n, \beta_n)$, where $t_i \in \mathbb{R}^+$ is a time point and $\beta_i \in B$. The ball $\beta \parallel P$ is the ball β translated by the process vector P ; that is, β translated by the concentration value for each dimension (species) in P .

Definition 11 (Trace and flow tube trace).

We have a function $\mathbf{trace} : \mathbb{P} \times \mathbb{R}^+ \times \mathbb{R}^+ \rightarrow \text{Trace}$, where $\mathbf{trace}(P, t, \rho)$ gives the trace of process P up to time t with resolution ρ using numerical simulation. We also have a function $\mathbf{tube} : B \times \mathbb{R}^+ \times \mathbb{R}^+ \rightarrow \text{Tube}$, where $\mathbf{tube}(\beta, t, \rho)$ gives the flow tube trace from the ball β for time t with resolution ρ ; each ball in the flow tube trace is given by $\Xi_{c(\beta), r(\beta)}(t_i)$ for each t_i up to t such that $t_{i+1} - t_i = \rho$.

The set \mathbb{B}_\perp is $\{\text{True}, \text{False}, \perp\}$ where \perp is of “uncertain” Boolean value. The set Signal_\perp is the set of finite length \mathbb{B}_\perp signals, signals with uncertainty, defined as follows:

Definition 12 (Finite length \mathbb{B}_\perp signal). A finite length \mathbb{B}_\perp signal s of length r is a function $s : [0, r) \rightarrow \mathbb{B}_\perp$. A finite length \mathbb{B}_\perp signal has finite variability and, therefore, may be represented by a finite interval covering. The interval covering, minimal covering, and definitions of consistency and refinement are the same as in Definition 7.

The set of positive intervals of s is $\mathcal{I}_s^+ = \{I \in \mathcal{I}_s : s(I) = \text{True}\}$, the set of uncertain intervals of s is $\mathcal{I}_s^\perp = \{I \in \mathcal{I}_s : s(I) = \perp\}$, and the set of negative intervals is $\mathcal{I}_s^- = \mathcal{I}_s \setminus (\mathcal{I}_s^+ \cup \mathcal{I}_s^\perp)$.

Definition 13 (\mathbb{B}_\perp signal combinators). The signal combinators apply the logical connectives (\neg, \wedge) and temporal modalities (\mathbf{F}, \mathbf{U}) to signals (s_ϕ, s_ψ) and are defined as follows:

$\neg s_\phi$

Negation is a simple negation of the signal such that $\mathcal{I}_{s_{\neg\phi}}^+ = \mathcal{I}_{s_\phi}^-$ and $\mathcal{I}_{s_{\neg\phi}}^\perp = \mathcal{I}_{s_\phi}^\perp$.

$s_\phi \wedge s_\psi$

For conjunction we first compute a refinement of the coverings $\mathcal{I}_\phi^R \prec \mathcal{I}_\phi$ and $\mathcal{I}_\psi^R \prec \mathcal{I}_\psi$ such that $\mathcal{I}_\phi^R = \mathcal{I}_\psi^R$ and is the sequence of intervals I_1^R, \dots, I_n^R . The conjunction is then computed interval-wise such that $s_{\phi \wedge \psi} = s_\phi \wedge s_\psi$ where $s_{\phi \wedge \psi} = \perp$ when either $s_\phi = \perp$ or $s_\psi = \perp$. The minimal covering $\mathcal{I}_{s_{\phi \wedge \psi}}$ is then computed by merging any adjacent intervals of the same value.

$\mathbf{F}_{[a,b]} s_\phi$

The temporal $\mathbf{F}_{[a,b]}$ modality is computed by back-shifting the positive and uncertain intervals. $\mathcal{I}_{\mathbf{F}_{[a,b]} s_\phi}^+$ is computed by taking each interval $I \in \mathcal{I}_\phi^+$ and computing its back-shifting $I \ominus [a, b] \cap \mathbb{R}^+$. In the same way $\mathcal{I}_{\mathbf{F}_{[a,b]} s_\phi}^\perp$ is computed and where any interval overlaps with an interval in $\mathcal{I}_{\mathbf{F}_{[a,b]} s_\phi}^+$ the overlapping portion of the interval is removed from $\mathcal{I}_{\mathbf{F}_{[a,b]} s_\phi}^\perp$. The minimal covering $\mathcal{I}_{\mathbf{F}_{[a,b]} s_\phi}$ is then computed by merging any adjacent intervals of the same Boolean value.

$s_\phi \mathbf{U}_{[a,b]} s_\psi$

The fundamental temporal $\mathbf{U}_{[a,b]}$ modality can be computed on the basis that $\phi \mathbf{U}_{[a,b]} \psi \iff \phi \wedge \mathbf{F}_{[a,b]}(\phi \wedge \psi)$ when s_ϕ is a unitary signal. A signal s is unitary if $\mathcal{I}_s^+ \cup \mathcal{I}_s^\perp$ is a singleton. So if s_ϕ is unitary and it has a value at t_1 and t_2 then it must hold that value for the whole interval $[t_1, t_2]$. For the case where s_ϕ is not unitary we can decompose it into a set of unitary signals $\{s_\phi^1, \dots, s_\phi^n\}$ and compute, for each $i \in [1, n]$:

$$s_{\phi \mathbf{U}_{[a,b]} \psi}^i = s_\phi^i \wedge \mathbf{F}_{[a,b]}(s_\phi^i \wedge s_\psi)$$

The signal is then recomposed to give:

$$s_{\phi \mathbf{U}_{[a,b]} \psi} = \bigvee_{i=1}^n s_{\phi \mathbf{U}_{[a,b]} \psi}^i$$

Our sensitive model checker for \mathcal{LBC} over $c\pi$ processes is defined by four mutually recursive functions. Function **sat** is the top-level function which computes whether or not a process satisfies an \mathcal{LBC} formula. Function **satB** computes whether a ball satisfies a formula: giving *True*, *False* or \perp according to whether every point of the ball satisfies the formula, none do, or only some. The function **signal** computes the formula satisfaction signal along a process trajectory. Finally, **signalT** computes a satisfaction signal along a flow tube trace. As with **satB**, this signal is three-valued according to whether the cross-section of the flow tube at each trace instant lies within, outside, or partially within the region satisfying the formula.

The algorithms for our sensitive model checker have two control parameters: a time resolution ρ as the step size for traces and flow tube traces; and θ , the ball radius within which we use the expansion function to extrapolate a flow tube.

Definition 14 (sat). *Function $\text{sat} : \mathbb{P} \times \Phi \rightarrow \mathbb{B}$ is computed recursively as follows:*

$$\begin{aligned} \text{sat}(P, \text{Atom}) &= \text{Atom} \in \text{Props}(P) \\ \text{sat}(P, \phi \wedge \psi) &= \text{sat}(P, \phi) \wedge \text{sat}(P, \psi) \\ \text{sat}(P, \neg\phi) &= \neg(\text{sat}(P, \phi)) \\ \text{sat}(P, (Q \triangleright \phi)) &= \text{sat}(P \parallel Q, \phi) \\ \text{sat}(P, \phi \mathbf{U}_{[a,b]} \psi) &= (\text{signal}(P, |\phi \mathbf{U}_{[a,b]} \psi|, \phi \mathbf{U}_{[a,b]} \psi))(0) . \end{aligned}$$

Satisfaction of a non-temporal formula is straightforward and can be determined directly from the initial conditions. Even a context formula, if its subformula is non-temporal, requires only a process composition and an inspection of initial conditions. However, for a temporal formula we compute the signal of its satisfaction over time; and then take the initial value of that signal. This leads us to the next function required.

Definition 15 (signal). *Function $\text{signal} : \mathbb{P} \times \mathbb{R}^+ \times \Phi \rightarrow \text{Signal}$ is computed recursively as follows, using the Boolean signal combinators from Definition 8:*

$$\begin{aligned} \text{signal}(P, t, \text{Atom}) &= \text{basicSignal}(P, t, \text{Atom}) \\ \text{signal}(P, t, \phi \wedge \psi) &= \text{signal}(P, t, \phi) \wedge \text{signal}(P, t, \psi) \\ \text{signal}(P, t, \neg\phi) &= \neg(\text{signal}(P, t, \phi)) \\ \text{signal}(P, t, \phi \mathbf{U}_{[a,b]} \psi) &= \text{signal}(P, t, \phi) \mathbf{U}_{[a,b]} \text{signal}(P, t, \psi) \\ \text{signal}(P, t, Q \triangleright \phi) &= \text{contextSignal}(P, t, Q, \phi) \end{aligned}$$

where **basicSignal** and **contextSignal** are defined below.

basicSignal : $\mathbb{P} \times \mathbb{R}^+ \times \text{Atomic} \rightarrow \text{Signal}$

Here **basicSignal**(P, t, Atom) gives the finite signal s such that for each $(t_i, c_i) \in \text{trace}(P, t, \rho)$: $s([t_i, t_i + \rho)) = (\text{Atom} \in \text{Props}(c_i))$.

contextSignal : $\mathbb{P} \times \mathbb{R}^+ \times \mathbb{P} \times \Phi \rightarrow \text{Signal}$

Function **contextSignal**(P, t, Q, ϕ) computes a signal s as follows. Take set X of all process states in the trace $\tau = \text{trace}(P, t, \rho)$, let β_X be the minimum bounding ball around X , and let I be the interval $[t_0, t_n + \rho)$ including time points t_0, \dots, t_n of τ . If $\text{satB}(\beta_X \parallel Q, \phi)$ is either True or False then $s(I) = \text{satB}(\beta_X \parallel Q, \phi)$ defines our signal s . If not, and τ contains only one time point (t, \mathbf{c}) , then take $s(I) = \text{sat}(\mathbf{c} \parallel Q, \phi)$. Finally, if satB gives \perp and τ contains multiple points, then bisect τ and repeat the procedure for each new τ , X and I .

The signal for an atomic proposition is computed directly from the simulation trace, interpolating between time points. The signal for non-atomic formulae other than context modalities is computed by using the appropriate signal combinators.

The difficult case is for context modalities. For this we compute a trace as for atomic propositions, translate it by the context Q , and then test a bounding ball around all points in the trace. If this is inconclusive we repeatedly refine until we find a set of balls which give a conclusive result. In the worst case this means checking individual points of the trace — as we did in our earlier methods. However, in any other case, we may save computation by checking a whole set of points together in a single ball.

That, however, requires a function to compute satisfaction across a ball.

Definition 16 (satB). The three-valued function $\text{satB} : B \times \Phi \rightarrow \mathbb{B}_\perp$ for satisfiability of a formula across a ball is computed recursively:

$$\begin{aligned} \text{satB}(\beta, \text{Atom}) &= \begin{cases} \text{True} & \beta \subseteq \{x \in \mathbb{P} \mid \text{Atom} \in \text{Props}(x)\} \\ \text{False} & \beta \cap \{x \in \mathbb{P} \mid \text{Atom} \in \text{Props}(x)\} = \emptyset \\ \perp & \text{otherwise} \end{cases} \\ \text{satB}(\beta, \phi \wedge \psi) &= \begin{cases} \text{True} & (\text{satB}(\beta, \phi) = \text{True}) \wedge (\text{satB}(\beta, \psi) = \text{True}) \\ \perp & (\text{satB}(\beta, \phi) = \perp) \vee (\text{satB}(\beta, \psi) = \perp) \\ \text{False} & \text{otherwise} \end{cases} \\ \text{satB}(\beta, \neg \phi) &= \begin{cases} \text{True} & \text{satB}(\beta, \phi) = \text{False} \\ \perp & \text{satB}(\beta, \phi) = \perp \\ \text{False} & \text{satB}(\beta, \phi) = \text{True} \end{cases} \\ \text{satB}(\beta, Q \triangleright \phi) &= \text{satB}(\beta \parallel Q, \phi) \\ \text{satB}(\beta, \phi \mathbf{U}_{[a,b]} \psi) &= \begin{cases} (\text{signalT}(\text{tube}(\beta, |\phi \mathbf{U}_{[a,b]} \psi|), \phi \mathbf{U}_{[a,b]} \psi))(0) & r(\beta) \leq \theta \\ \perp & \text{otherwise} \end{cases} \end{aligned}$$

For atomic propositions, logical combinations, and the context modality, we need only determine whether a ball lies entirely inside or outside the region defined by the formula; with some refinement for combinations with the mixed value \perp . For temporal formulae, if the ball has a radius too large for reliable

extrapolation by the expansion function, then we return \perp . Note, though, that where **satB** has been called from the within the loop of **contextSignal** this will immediately trigger bisection and further calls to **satB** over a smaller region. Finally, to compute validity of a temporal formula over a ball of radius θ or less, we compute the signal over a flow tube trajectory and take its initial value.

That, of course, requires that we compute a signal over the tube trace.

Definition 17 (signalT). *The function $\text{signalT} : \text{Tube} \times \Phi \rightarrow \text{Signal}_\perp$ is defined recursively as follows:*

$$\begin{aligned} \text{signalT}(T, \text{Atom}) &= \text{basicSignalT}(T, \text{Atom}) \\ \text{signalT}(T, \phi \wedge \psi) &= \text{signalT}(T, \phi) \wedge \text{signalT}(T, \psi) \\ \text{signalT}(T, \neg\phi) &= \neg(\text{signalT}(T, \phi)) \\ \text{signalT}(T, \phi \mathbf{U}_{[a,b]} \psi) &= \text{signalT}(T, \phi) \mathbf{U}_{[a,b]} \text{signalT}(T, \psi) \\ \text{signalT}(T, Q \triangleright \phi) &= \text{contextSignalT}(T, Q, \phi) . \end{aligned}$$

*This uses the operations on three-valued signals from Definition 13 and auxiliary functions **basicSignalT** and **contextSignalT**.*

basicSignalT: $\text{Tube} \times \text{Atomic} \rightarrow \text{Signal}_\perp$

For this $\text{basicSignalT}(T, \text{Atom})$ is the finite three-valued signal s where for each $(t_i, \beta_i) \in T$ we have $s([t_i, t_i + \rho)) = \text{satB}(\beta_i, \text{Atom})$.

contextSignalT: $\text{Tube} \times \mathbb{P} \times \Phi \rightarrow \text{Signal}_\perp$

Function $\text{contextSignalT}(T, Q, \phi)$ computes a three-valued signal as follows. Let W be the set of all balls in the tube trace T , let β_W be the minimum bounding ball around W , and let I be the interval $[t_0, t_n + \rho)$ including time points t_0, \dots, t_n of T . If $\text{satB}(\beta_W \parallel Q, \phi)$ is either True or False then we have our signal defined by $s(I) = \text{satB}(\beta_W \parallel Q, \phi)$. Otherwise, if the tube trace T contains only a single time point (t, β) then $s(I) = \text{satB}(\beta \parallel Q, \phi)$. Finally, if satB gives \perp and T contains multiple time points, then bisect T and repeat the procedure for each new tube T , set of balls W and interval I .

Computation of a three-valued signal for a flow tube is very similar to that for the boolean signal of a trajectory, except that we deal in balls rather than points. However, where in **contextSignal** we group points into a ball to send to **satB**, here we group balls into their bounding ball, and still need only call **satB** — closing off the mutually recursive definition of our four functions.

The key advance here is that by using sensitivity calculations to drive the expansion function, we may replace multiple calls to a numerical solver to repeatedly compute traces with a single call to compute a flow tube. The full power of this is engaged when we have alternation of temporal and contextual modalities, and may need to calculate trajectories starting at every point along a spinal translated trajectory. The sensitive model-checker groups trajectories starting within θ of each other together into a single flow tube and a single signal. In the best case, we may tremendously reduce the number of calls to the numerical solver and the number of signals to compute; in the worst case, subdivision leads us to the previous algorithm of a trace at every point.

This reduction in computation addresses the second problem of §4.2. Our algorithm also begins to address the first problem there, of intermediate values between discrete points on a trace. When `satB` evaluates a formula of \mathcal{LBC} across a ball including several points of a trajectory, it also does so for all values between them. This extends to flow tubes, too, surrounding all trajectories that start in their initial ball. This considerably extends the earlier, purely pointwise, algorithm. However, it is not necessarily complete as a very volatile trajectory may conceivably range outside this enclosing ball in between trace points.

6 Conclusion

In this paper we have shown ways to address two problems associated with using numerical ODE solvers for model-checking contextual properties of continuous dynamical systems. We use methods from signal logics and from sensitivity analysis.

The first problem was that of computational cost, arising from checking a very large number of time-points along execution traces from the numerical solver. The use of signal-logic methods reduces the overhead in working with large traces, and for the context modality in particular we have identified ways to use sensitivity analysis in numerical solvers to check a covering collection of small regions rather than many individual points. This aims to substantially reduce the number of subsequent calls to the solver to check nested spatial modalities.

The second problem was related to the correctness of the model checker when working with discrete-time traces. By computing with flow tubes around trajectories that cover a computed set of initial conditions, we have a stronger guarantee that we have not missed additional trajectories from intermediate time points. However, it is still possible to miss initial conditions that lie outwith the bounding ball of the discrete time points on a trace. As would be expected, the correctness of this model-checking also remains dependent on the accuracy of the numerical solver.

We have done some implementation, validating these approaches. In particular, we have a full implementation of the signal-logic techniques of §4 in our existing \mathcal{LBC} -checker, and can confirm a substantial performance improvement from the highly-compressed signal representation of traces. We have not yet implemented the methods using expansion functions and sensitivity analysis of §5.

Acknowledgements

We are very grateful to Luca Bortolussi for suggesting the use of signals, and introducing us to the work of Maler et al.; and to others in the Edinburgh PEPA group. This work was supported by the Laboratory for Foundations of Computer Science, and by the Engineering and Physical Sciences Research Council under grant number EP/P50550X/1.

References

1. ALUR, R., FEDER, T., AND HENZINGER, T. The benefits of relaxing punctuality. *Journal of the ACM (JACM)* 43, 1 (1996), 116–146.
2. ANTONIOTTI, M., POLICRITI, A., UGEL, N., AND MISHRA, B. Model building and model checking for biochemical processes. *Cell biochemistry and biophysics* 38, 3 (Jan. 2003), 271–86.
3. BANKS, C. J., AND STARK, I. A Logic of Behaviour in Context. *Information and Computation* (2014), in press.
4. CAIRES, L., AND LOZES, E. Elimination of quantifiers and undecidability in spatial logics for concurrency. *Theoretical computer science* 358, 2-3 (2006), 293–314.
5. CALZONE, L., CHABRIER-RIVIER, N., FAGES, F., AND SOLIMAN, S. Machine learning biochemical networks from temporal logic properties. *Transactions on Computational Systems Biology VI* (2006), 68–94.
6. CARDELLI, L., AND GORDON, A. Anytime, anywhere: Modal logics for mobile ambients. In *Proceedings of the 27th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (2000), ACM, pp. 365–377.
7. DE NICOLA, R. D., AND LORETI, M. MoMo: A modal logic for reasoning about mobility. In *Formal Methods for Components and Objects* (2005), Springer, pp. 95–119.
8. DONZÉ, A., AND MALER, O. Systematic simulation using sensitivity analysis. *Hybrid Systems: Computation and Control* (2007), 174–189.
9. FAGES, F., AND RIZK, A. On temporal logic constraint solving for analyzing numerical data time series. *Theoretical Computer Science* 408, 1 (Nov. 2008), 55–65.
10. KWIATKOWSKI, M. *A formal computational framework for the study of molecular evolution*. PhD thesis, University of Edinburgh, 2010.
11. KWIATKOWSKI, M., AND STARK, I. The continuous π -calculus: A process algebra for biochemical modelling. In *Computational Methods in Systems Biology* (2008), Springer, pp. 103–122.
12. MALER, O., AND NICKOVIC, D. Monitoring temporal properties of continuous signals. In *Formal Techniques, Modelling and Analysis of Timed and Fault-Tolerant Systems* (2004), Springer, pp. 152–166.
13. MILNER, R. *Communicating and mobile systems: the pi-calculus*. Cambridge University Press, 1999.
14. NICKOVIC, D., AND MALER, O. AMT: A property-based monitoring tool for analog systems. *Formal Modeling and Analysis of Timed Systems* (2007).
15. REGEV, A., AND SHAPIRO, E. Cells as computation. *Nature* 419 (2002), 343.
16. REGEV, A., SILVERMAN, W., AND SHAPIRO, E. Representation and simulation of biochemical processes using the pi-calculus process algebra. In *Pacific Symposium on Biocomputing* (2001), vol. 6, World Scientific Pub Co Inc, pp. 459–470.